arXiv:1404.4388v1 [math.OC] 16 Apr 2014

Partially Observed, Multi-objective Markov Games

Yanling Chang, Alan L. Erera, and Chelsea C. White III

School of Industrial and Systems Engineering
Georgia Institute of Technology

Abstract

The intent of this research is to generate a set of non-dominated policies from which one of two agents (the leader) can select a most preferred policy to control a dynamic system that is also affected by the control decisions of the other agent (the follower). The problem is described by an infinite horizon, partially observed Markov game (POMG). The actions of the agents are selected simultaneously at each decision epoch. At each decision epoch, each agent knows: its past and present states, its past actions, and noise corrupted observations of the other agent's past and present states. The actions of each agent are determined by the agent's policy, which selects actions at each decision epoch based on these data.

The leader considers multiple objectives in selecting its policy. The follower considers a single objective in selecting its policy with complete knowledge of and in response to the policy selected by the leader. This leader-follower assumption allows the POMG to be transformed into a specially structured, partially observed Markov decision process (POMDP). This POMDP is used to determine the follower's best response policy. A multi-objective genetic algorithm (MOGA) is used to create the next generation of leader policies based on the fitness measures of each leader policy in the current generation. Computing a fitness measure for a leader policy requires a value determination calculation, given the leader policy and the follower's best response policy. The policies from which the leader can select a most preferred policy are the non-dominated policies of the final generation of leader policies created by the MOGA. An example is presented that illustrates how these results can be

used to support a manager of a liquid egg production process (the leader) in selecting a sequence of actions to best control this process over time, given that there is an attacker (the follower) who seeks to contaminate the liquid egg production process with a chemical or biological toxin.

1 Introduction

The intent of this research is to provide decision support to a *leader*, an agent that wishes to determine a policy that will select actions to best control a sequential stochastic system over an infinite planning horizon, given that there is a *follower*, an agent that also would like to exert control over the system for its own purposes. We assume:

- At each decision epoch, each agent knows: its past and present states, its past actions, and noise corrupted observations of the other agent's past and present states.
- The leader's and follower's actions are selected simultaneously at each decision epoch.
- Each agent's policy selects actions based on data currently available to the agent.
- The follower knows the leader's policy and determines a response policy that is optimal with respect to the follower's objective.
- The leader considers multiple objectives in selecting its policy.

Our objective is to find a set of non-dominated policies from which the leader can select its most preferred policy. Such information can serve as input to a decision support system that, for example, is based on a deterministic version of multi-attribute utility theory (Keeney & Raiffa, 1993; Holloway & White, 2008). In this context, the results presented in this paper generate options (i.e., policies) for consideration by (1) creating multiple generations of policies and eliminating all but the non-dominated set of policies from the last generation and (2) determining value scores for each of the policies in this non-dominated set.

We remark that the assumption that the follower knows the policy of the leader is a conservative assumption from the perspective of the leader and could unrealistically bias

the game to the advantage of the follower. However, this bias is mollified by the fact that the leader and follower do not share the same data at each decision epoch, and hence the follower can only infer what action the leader will actually take. This assumption is also reasonable for many applications. For example, if the leader is a large governmental agency or corporation and the follower is an individual or group intent on attacking the leader, then it may be reasonable to assume that the follower will know more about the leader than the leader will know about the follower. Further, the assumption that the follower knows the policy of the leader allows us to transform the Markov game that we use to model leader-follower interaction into a model of sequential decision making under uncertainty and hence to take advantage of this computationally useful transformation.

The motivating application of this research is the operation of a liquid egg production facility in order to maximize the supply chain's productivity while minimizing its vulnerability to the intentional insertion of a biological or chemical toxin into the food production and distribution system (see Manning, 2005; O'Ryan, 1996; Sobel, 2002). For this application, we assume the leader manages the production facility and is trying to balance two objectives: (1) maximize productivity and (2) minimize vulnerability. See Mohtadi & Murshid (2009) for background information about this application area. Although initially developed to model this application, we remark that the decision support process to be presented can model a particularly broad class of sequential game applications, when all agents are intelligent and adaptive.

Many of the methodological characteristics of the decision support model presented in this paper have been considered elsewhere in the decision, risk, and reliability analysis literatures. Models of intelligent agents or adversaries are examined by Cardoso & Diniz (2009). The single-period leader-follower game has been widely used to analyze the strategic interactions between two intelligent and adaptive agents. For example, Cavusoglu et al. (2013) have studied the impacts of passenger profiling on airport security operations. Bakir (2011) analyzed resource allocation for cargo container transportation

security. Other applications of the single-period leader-follower game are presented in Bier (2007, 2008) and Zhuang & Bier (2007). Multi-period games have been considered by Wang & Bier(2011), who examined a two-period leader-follower repeated game, and by Hausken & Zhuang (2011), who studied a multi-period game with myopic agents. Application of completely observable stochastic game to overseas cargo container security can be found in Bakir & Kardes (2009) to capture the state dynamics over time. Models that consider incomplete or uncertain information are presented and analyzed by McLay (2012), Rothschild (2012) and Wang & Bier (2011).

Each of above methodological characteristics is intended to enhance the realism of the respective model. Our model extends the existing literature on sequential games by explicitly considering the multi-period interaction of two non-myopic agents, each of whom adjust its decisions according to the other agent's decisions over an infinite planning horizon. Furthermore, our model considers the case where neither agent has complete information about the other agent. By combining these characteristics into a single model, as has been done in this paper, we believe that the modeling realism of the resulting model has been further enhanced. However, and not surprisingly, additional model realism has resulted in increased computational challenges. Dealing with these challenges is the focus of much of this paper.

Our approach to decision support is described as follows. We begin with an initial (i.e., first generation) set of possible leader policies. We then use a multi-objective genetic algorithm (MOGA) to create successive generations of leader policies. Presumably, the next generation of leader policies contains, in some sense, higher quality policies than the current generation. We then determine the non-dominated set of the last generation of leader policies determined and present this set to the leader. The leader can then select the most preferred policy from this set for implementation.

Mimicking the process of natural evolution, the MOGA creates the next generation of policies from the current set based on the fitness measures of each of the policies in

the current set. Each fitness measure is related to an objective of the leader. We model the interaction between the leader and follower as a partially observed Markov game (POMG). Our POMG is a version of the partially observed stochastic game (POSG) where the state dynamics possess the Markovian property. We assume that the follower is aware of the policy that the leader has selected and makes use of this fact in constructing the follower's policy. Thus at the policy level, the game is a leader-follower (Stackelberg) game. This assumption allows the POMG to be converted into a partially observed Markov decision process (POMDP). The fitness measures for each leader policy in the current generation are needed by the MOGA to create the next generation of leader policies from the current generation. These fitness measure are computed by a value determination procedure, given the leader policy and the follower response policy.

The paper is organized as follows. We review the pertinent literature associated with the MOGA, POSG, and POMDP in Section 2. In Section 3, we describe the MOGA in more detail, show how the fitness measures are computed using the POMG and the POMDP, present equilibrium conditions, and discuss the value of information. The numerical evaluation in Section 4 applies this decision support procedure to a simplified liquid eggs supply chain security problem and analyses the value of information for the agents. Section 5 summarizes research results and discusses future research directions.

2 Literature Review

The research presented in this paper combines and extends results associated with POMDP, POSG, and MOGA. We now review the pertinent literature in these three areas of research.

2.1 The partially observed Markov decision process

The POMDP is a model of sequential decision making under uncertainty that takes into consideration noise corrupted and/or costly observations of the state of the system under control. Relative to the completely observed Markov decision process (i.e., the MDP; see Puterman, 1994), the POMDP represents a more general but significantly more computationally challenging model. In seminal research, Smallwood and Sondik (1973) and Sondik (1978) showed that under robust conditions the optimal cost function for both the finite horizon and infinite horizon expected total discounted cost criterion POMDPs is piecewise linear and convex, and presented successive approximations approaches for solving these POMDPs that exploited this structure. Zhang (2010) revisited these structural results and convergence properties by exploiting the dual relationship between hyperplanes and points in the POMDP and related the solution of the POMDP to the Minkowski sum problem in computational geometry. Monahan (1982), Eagle (1984), and White and Scherer (1989) presented improved algorithms based on the structural results. Detailed descriptions of other exact algorithms can be found in Cheng (1988), Littman (1994b), Cassandra (1994a), Cassandra, Littman and Zhang (1997), Feng and Zilberstein (2004), Lin and White (1998, 2004) and Naser-Moghadasi (2012). Surveys of related solution techniques and complexity analyses for the POMDP can be found in Monahan (1982), Lovejoy (1991), White (1991), Cassandra (1994b) and Poupart (2005).

In the development of approximate solution techniques for POMDPs, point-based value iteration (PBVI) is presented and analysed in Pineau (2003) and Shani (2012). Platzman (1977, 1980), White and Scherer (1994), Littman (1994a), Hauskrecht (1997), Hansen (1998a, 1998b), Poupart (2005) examined finite memory policy and finite-state controllers. Varakantham (2007) and Poupart (2011) focused on calculating bounds on optimal POMDP solutions in order to evaluate the quality of approximate solutions. Surveys of approximation methods for POMDPs can be found in Hauskrecht (2000), Aberdeen (2003), and Yu (2007). Approximate algorithms have proved useful for large-

scale problems (Hoey, 2010; Thomson and Young, 2010).

2.2 The partially observable stochastic game

The stochastic game introduced by Shapley (1953) represents a multi-agent planning problem in a stochastic environment. In this setting, each player considers the consequences of its own action and the actions that its opponents or teammates may take. See Raghavan and Filar (1991), Filar (1997) and Ummels (2010) for details.

The POSG is a new, relatively unexamined generalization of the stochastic game, where the states of the game are not precisely observed by the players and all players make decisions based on these partial observations. Although POSGs provide a robust framework for multi-agent planning, Bernstein (2002) showed that POSGs are computationally intractable when problem size grows. Rabinovich (2003) has shown that even epsilon-optimal approximations are NP-hard. As a result, POSGs with special structures that enhance computational tractability are of considerable interest. Koller (1994) provided an efficient algorithm for a two-player POSG with tree-like structure. McEneaney (2004) focused on a game where only one player has imperfect information. Ghosh (2004), Oliehoek (2005), and Bopardikar (2011) studied a zero-sum version of the POSG. Emery-Montemerlo (2004) approximated POSGs with common payoffs by a series of Bayesian games. A cooperative version of the POSG, called a decentralized partially observable Markov decision process (DEC-POMDP), has been studied by Becker (2004), Bernstein (2005), Seuken (2007), and Oliehoek (2008). A survey of the DEC-POMDP can be found in Oliehoek (2012).

It is often the case for real-world planning problems that the players' payoffs are neither completely aligned with others nor directly opposed. Hespanha and Prandini (2001) proved the existence of Nash equilibrium in a two-player finite-horizon POSG. Hansen (2004) developed a dynamic program for general POSGs by pruning very weakly dominated strategies and then showed that this dynamic programming approach can

achieve optimality for cooperative settings. However, this approach is computationally infeasible for all but the smallest problems. Kumar and Zilberstein (2009) developed an approximate solution procedure for the POSG based on Hansen's work. Interactive POMDPS addressed in Gmytrasiewicz and Doshi (2005) demonstrated another framework for multi-agent planning.

2.3 The multi-objective genetic algorithm

Genetic algorithms, introduced by Holland (1975), are adaptive heuristic search techniques that mimic the process of natural evolution. A genetic algorithm represents each feasible problem solution in a population of solutions as a genome or chromosome, and begins with an initial population of feasible solutions. Solutions having high measures of fitness are preferably selected during each generation to produce the next generation of solutions having improved fitness measures by applying genetic (e.g., mutation and crossover) operators. After a number of generations, the population presumably evolves to optimal or near-optimal solutions. Goldberg (1989), Forrest (1993) and Srinivas(1994) present surveys of genetic algorithms and the theories.

Multi-objective genetic algorithms (MOGA) are designed for the simultaneous optimization of multiple, often competing objectives. Usually the optimal solutions are a set of points, called the Pareto-optimal set, in the sense that no improvement can be made in any objective without sacrificing the other objectives. MOGA pushes the Pareto frontier towards the ideal optimal set of solutions as the algorithm proceeds. MOGA algorithms include: the vector evaluated GA (VEGA) (Schaffer, 1985), the Niched Pareto GA (NPGA) (Horn, 1994), the Pareto Envelope-based Selection Algorithms (PESA) (Corne, 2000) and the Fast Non-dominated sorting GA (NSGA-II) (Deb, 2002). Surveys are presented in Coello (2000) and Konak (2006). MOGAs have been widely applied in optimization and decision making problems (see Ponnambalam, 2001; Deb, 2001; Ombuki, 2006; Lin, 2008; Bowman, 2010; Yildirim, 2012).

3 Model and Analysis

We first present the POMG model in Section 3.1. In order to determine an optimal response policy for the follower, the POMDP is constructed by combining the POMG model with any leader's policy. The resulting POMDP is presented and examined in Section 3.2. For computational reasons, we require that the leader and follower policies be finite memory policies. However, the POMDP constructs a perfect memory follower policy. In Section 3.3, we present an approach for determining a finite-memory approximation of a perfect memory policy. In order for the MOGA to determine the next generation of leader policies, fitness measures must be calculated for each policy in the current generation of leader policies. Each fitness measure is associated with an objective of the leader. In Section 3.4 we present an approach for determining the fitness measures, for any given leader policy and follower policy. Section 3.5 shows how to use MOGA to generate a non-dominate set of leader policies. Section 3.6 and 3.7 address equilibria and the value of information, respectively.

3.1 Partially Observed Markov Game

The partially observed Markov game (POMG) serves as the modeling basis of our decision support system design. The POMG is comprised of:

Decision epochs: Let $\{0, 1, \dots\}$ be the set of all decision epochs when both agents select actions simultaneously. Thus, the problem horizon is countable and infinite.

State spaces: Let S^L and S^F be the state spaces of the leader and the follower, respectively. Both spaces are epoch-invariant and finite. At decision epoch t, let $s^L(t)$ be the leader's state, $s^F(t)$ be the follower's state, and denote $s(t) = \{s^L(t), s^F(t)\}$.

Action spaces: Let A^L and A^F be the epoch-invariant action spaces of the leader and the follower, both of which are finite. At decision epoch t, let $a^L(t)$ be the leader's action, $a^F(t)$ be the follower's action, and denote $a(t) = \{a^L(t), a^F(t)\}$.

Observation spaces: Let Z^L and Z^F be the observation spaces of the leader and the follower, both of which are epoch-invariant and finite. At decision epoch t, let $z^F(t)$ be the follower's observation of the leader's state, $z^L(t)$ the leader's observation of the follower's state, and denote $z(t) = \{z^L(t), z^F(t)\}$.

Systems dynamics: We assume the epoch-invariant probability $P(z(t+1), s(t+1)|s(t), a(t))$ is given. Note that

$$P(z(t+1), s(t+1)|s(t), a(t)) = P(z(t+1)|s(t+1), s(t), a(t))P(s(t+1)|s(t), a(t)),$$

where $P(z(t+1)|s(t+1), s(t), a(t))$ is referred to as the state observation probability and $P(s(t+1)|s(t), a(t))$ is referred to as the state transition probability.

Information patterns: The information pattern for agent k describes what agent k knows and when agent k knows it. Let

- $Z^k(t) = \{z^k(t), z^k(t–1), \dots\}$ and $Z^k(t, \tau) = \{z^k(t), z^k(t–1), \dots, z^k(t–\tau+1)\}$
- $S^k(t) = \{s^k(t), s^k(t–1), \dots\}$ and $S^k(t, \tau) = \{s^k(t), s^k(t–1), \dots, s^k(t–\tau+1)\}$
- $\mathcal{A}^k(t) = \{a^k(t-1), a^k(t–2), \dots\}$ and $\mathcal{A}^k(t, \tau) = \{a^k(t-1), \dots, a^k(t–\tau)\}$
- $\mathcal{I}^k(t) = \{\mathcal{Z}^k(t), \mathcal{S}^k(t), \mathcal{A}^k(t)\}$
- $\mathcal{I}^k(t, \tau) = \{\mathcal{Z}^k(t, \tau), \mathcal{S}^k(t, \tau), \mathcal{A}^k(t, \tau)\}$

We assume that agent k chooses $a^k(t)$ on the basis of $\mathcal{I}^k(t)$, if agent k has perfect memory, or on the basis of $\mathcal{I}^k(t, \tau)$, if agent k has finite memory. Note that $\mathcal{I}^k(t) = \{\mathcal{I}^k(t, \tau), \mathcal{I}^k(t-\tau)\}$.

Single Period Cost and Criteria: Let $c^F(s(t), a(t))$ be the decision epoch invariant single period cost accrued by the follower at epoch t, given $s(t)$ and $a(t)$, and let $c_i^L(s(t), a(t))$ be the decision epoch invariant single period cost accrued by the leader with respect to criterion i at epoch t, given $s(t)$ and $a(t)$. The criteria under consideration are the concomitant expected total discounted costs over the infinite horizon.

Policies: A policy π^k for agent k is a mapping from what agent k knows at epoch t, either $\mathcal{I}^k(t)$ or $\mathcal{I}^k(t,\tau)$, into its set of available actions, A^k. Policies can be random and hence described by conditional probabilities. We restrict our interest to stationary policies. Stationary policies tend to be easy to implement and in many cases, e.g., the determination of optimal follower response policies for a broad class of scalar criteria, sufficiently rich to contain an optimal policy.

Objectives: The follower's objective is to select a stationary policy that minimizes its criterion. The leader's objective is to optimize all criteria under consideration in some balanced manner, with this balance being determined by the leader. Our objective is to provide the leader with a non-dominated set of policies from which to choose a single policy for implementation.

3.2 Determination of a Best Response Policy $\bar{\pi}^F$, Given a leader policy π^L

Let $\bar{\pi}^F$ be the *perfect* memory best response policy to the leader policy π^L. We assume that the leader policy $\pi^L = \{P(a^L(t)|\mathcal{I}^L(t,\tau))\}$ and that the follower knows π^L. We also assume that the follower knows $\mathcal{I}^F(t)$. (We will restrict the follower's information pattern to $\mathcal{I}^F(t,\tau)$ below; however, for the moment it will be convenient to assume that the follower has perfect memory.) We remark that although the follower knows the leader's policy, because the information patterns of the agents are in general different, the follower can only infer what action the leader will actually take.

Let $v^F(\pi^L;\mathcal{I}^F(t))$ be the follower's optimal criterion value, given $\mathcal{I}^F(t)$ and π^L. For notational simplicity, assume the dependence of π^L is implicit; hence, $v^F(\mathcal{I}^F(t)) = v^F(\pi^L;\mathcal{I}^F(t))$. Then, according to results in (Puterman, 1994; Chapter 6), v^F uniquely satisfies

$$v^F = H^F v^F \tag{1}$$

where for any v,

$$[H^F v](\mathcal{I}^F(t)) = \min_{a^F(t)} h^F(\mathcal{I}^F(t), a^F(t), v),$$

$$h^F(\mathcal{I}^F(t), a^F(t), v) = E\{c^F(s(t), a(t)) + \beta v(\mathcal{I}^F(t+1))|\mathcal{I}^F(t), a^F(t)\},$$

and where E is the expectation operator and the minimum is over all $a^F(t)$. We now state our first result.

Proposition 1. *Assume π^L is given. Then for each $s^F(t)$ there is an at most countable set of arrays $\Gamma^*(s^F(t))$ such that:*

$$v^F(\mathcal{I}^F(t)) = \min\{\sum \gamma(\mathcal{I}^L(t,\tau))P(\mathcal{I}^L(t,\tau)|\mathcal{I}^F(t)) : \gamma \in \Gamma^*(s^F(t))\},$$

where the sum is over all $\mathcal{I}^L(t,\tau)$.

Proof. Assume v and Γ are such that

$$v(\mathcal{I}^F(t)) = \min\{\sum \gamma(\mathcal{I}^L(t,\tau))P(\mathcal{I}^L(t,\tau)|\mathcal{I}^F(t)) : \gamma \in \Gamma(s^F(t))\},$$

where the sum is over all $I^L(t,\tau)$. Then straightforward analysis, following arguments in (Smallwood and Sondik, 1973), shows that

$$h^F(\mathcal{I}^F(t), a^F(t), v) = \min\{\sum \gamma'(\mathcal{I}^L(t,\tau))P(\mathcal{I}^L(t,\tau)|\mathcal{I}^F(t)) : \gamma' \in \Gamma'(s^F(t), a^F(t))\},$$

where the sum is over all $\mathcal{I}^L(t,\tau)$, where if $\gamma' \in \Gamma'(s^F(t), a^F(t))$ then γ' is of the form

$$\gamma'(\mathcal{I}^L(t,\tau)) = \sum_{a^L(t)} P(a^L(t)|\mathcal{I}^L(t,\tau))[c^F(s(t), a(t))$$
$$+ \beta \sum_{s(t+1)} \sum_{z(t+1)} \gamma^{i,j}(z^L(t+1), s^L(t+1), a^L(t), \mathcal{I}^L(t,\tau-1))P(z(t+1), s(t+1)|s(t), a(t))],$$

where $\gamma^{i,j}$ can be any element in $\Gamma(s^F(t+1))$ for each $s^F(t+1) = i$ and $z^F(t+1) = j$.

Then,

$$[H^F v](\mathcal{I}^F(t)) = \min\{ \sum_{\mathcal{I}^L(t,\tau)} \gamma''(\mathcal{I}^L(t,\tau)) P(\mathcal{I}^L(t,\tau)|I^F(t)) : \gamma'' \in \Gamma''(s^F(t))\},$$

where $\Gamma''(s^F(t)) = \cup_{a^F(t)} \Gamma'(s^F(t), a^F(t))$.

The operator H^F is a contraction operator on the Banach space comprised of all functions mapping $\mathcal{I}^F(t)$ into the real line, having as its norm the supremum norm, and as a result, the sequence $\{v^n\}$, where $v^{n+1} = H^F v^n$, converges to v^F for any given v^0. The above result indicates that H^F preserves piecewise linearity and concavity and in the limit preserves concavity. □

We remark that $\Gamma''(s^F(t))$ usually contains many redundant vectors, where γ is redundant if $[H^F v](\mathcal{I}^F(t))$ is strickly less than $\sum_{\mathcal{I}^L(t,\tau)} \gamma(\mathcal{I}^L(t,\tau)) P(\mathcal{I}^L(t,\tau)|\mathcal{I}^F(t))$ for all $\{P(\mathcal{I}^L(t,\tau)|\mathcal{I}^F(t))\}$. From both storage and computational perspectives, there is a value to keep the cardinality of $\Gamma''(s^F(t))$ as small as possible. Let the operator **PURGE** be such that $\Gamma'''(s^F(t)) = \mathbf{PURGE}(\Gamma''(s^F(t)))$ is the subset of $\Gamma''(s^F(t))$ having the smallest cardinality that satisfies

$$\begin{aligned}
&\min\{ \sum_{I^L(t,\tau)} \gamma''(\mathcal{I}^L(t,\tau) P(\mathcal{I}^L(t,\tau)|\mathcal{I}^F(t))) : \gamma'' \in \Gamma''(s^F(t+1))\} \\
&= \min\{ \sum_{I^L(t,\tau)} \gamma'''(\mathcal{I}^L(t,\tau) P(\mathcal{I}^L(t,\tau)|\mathcal{I}^F(t))) : \gamma'' \in \Gamma'''(s^F(t+1))\}
\end{aligned}$$

for all $\{P(\mathcal{I}^L(t,\tau)|\mathcal{I}^F(t))\}$. Related discussion about the necessity and the existence of the **PURGE** operator can be found in Lin and White (1998).

With regard to the implications of Proposition 1 and results in (Puterman, 1994; Chapter 6), v^F and hence an optimal policy can depend on $\mathcal{I}^F(t)$ only through $(s^F(t), y^F(t))$, where $y^F(t) = \{P(\mathcal{I}^L(t,\tau)|\mathcal{I}^F(t)) : \text{ for all } \mathcal{I}^L(t,\tau)\}$. Hence, $(s^F(t), y^F(t))$ is a sufficient statistic. Further, v^F is concave in $y^F(t)$. Additionally, if $\gamma^*(s^F(t))$ is a finite set of ar-

rays for all $s^F(t)$, then v^F is piecewise linear. Note that the dimension of $(s^F(t), y^F(t))$ is finite and t-invariant. Note further that the finite dimensionality of $y^F(t)$ follows directly from the finite-memory assumption imposed on π^L. Thus, assuming $\gamma^*(s^F(t))$ is a finite set of arrays and v^F is described in terms of $(s^F(t), y^F(t))$, v^F has a finite representation. We remark that the cardinality of $\Gamma'(s^F(t))$ can be substantially larger than the cardinality of $\Gamma(s^F(t))$, where both $\Gamma(s^F(t))$ and $\Gamma'(s^F(t))$ are defined in the proof of Proposition 1. Techniques for reducing the cardinality of $\Gamma'(s^F(t))$ can be found in White (1991).

3.3 A Finite-Memory Approximation to $\bar{\pi}^F$

As noted above, an optimal policy $\bar{\pi}^F$ for the follower that achieves the minimum in Equation 1 depends on $\mathfrak{I}^F(t)$ and hence $\bar{\pi}^F$ is a perfect-memory policy. In order to insure that the leader criteria have finite representation, the follower policy must also be a finite-memory policy. We determine a finite-memory (approximate) policy from a given perfect memory policy as follows. We note that $\{(\mathfrak{I}^F(t,\tau), y^F(t-\tau)), t = 1, 2, \dots\}$ is also a sufficient statistic for this problem, with $y^F(t-\tau)$ representing the influence of data determined up through epoch t–τ. By (arbitrarily) assuming a uniform distribution over $y^F(t-\tau)$, we determine probabilities of the form $P(a^F(t)|\mathfrak{I}^F(t,\tau))$. Our numerical analyses indicate that these finite memory approximations of optimal perfect memory policies can be remarkably accurate, even for small τ. Identifying and analyzing other approaches for determining a finite-memory policy from a perfect memory policy is a topic for future consideration. In the following context, we use π^F to denote a *finite* memory best response policy and $\bar{\pi}^F$ to represent a *perfect* memory best response policy of the follower.

We remark that an alternative approach for directly determining a finite-memory follower policy in response to a given leader policy involves determining v^F as a function of $\mathfrak{I}^F(t,\tau)$, rather than as a function of $\mathfrak{I}^F(t)$. Whether or not such an approach could

be useful is a topic of future research; see Platzman (1977) and White (1994) for related discussion.

3.4 Fitness Measure Determination

Let $v_i^L(\pi^L, \pi^F; \mathcal{I}^L(t))$ be the criterion value for the leader's i^{th} criterion, given $\mathcal{I}^L(t)$, a leader policy $\pi^L = \{P(a^L(t)|\mathcal{I}^L(t,\tau))\}$, and a follower policy $\pi^F = \{P(a^F(t)|\mathcal{I}^F(t,\tau))\}$. For notational simplicity, we assume that the dependence of $v_i^L(\pi^L, \pi^F; \mathcal{I}^L(t))$ on (π^L, π^F) is implicit; hence, $v_i^L(\mathcal{I}^L(t)) = v_i^L(\pi^L, \pi^F; \mathcal{I}^L(t))$. Then according to results in (Puterman, 1994; Chapter 6), v_i^L uniquely satisfies

$$v_i^L(\mathcal{I}^L(t)) = h_i^L(\mathcal{I}^L(t), v_i^L)$$

for all $\mathcal{I}^L(t)$, where

$$h_i^L(\mathcal{I}^L(t), v) = E\{c_i^L(s(t), a(t)) + \beta v(\mathcal{I}^L(t+1))|\mathcal{I}^L(t)\}.$$

We now show that $v_i^L(\mathcal{I}^L(t))$ is dependent on $\mathcal{I}^L(t)$ only though $(\mathcal{I}^L(t,\tau), y^L(t))$, where the array $y^L(t) = \{P(I^F(t,\tau)|\mathcal{I}^L(t)) : \text{ for all } \mathcal{I}^F(t,\tau)\}$. Thus, $(\mathcal{I}^L(t,\tau), y^L(t))$ is a sufficient statistic.

Proposition 2. *Assume (π^L, π^F) are given as two finite-memory policies. Then, there is a function g_i^* such that*

$$v_i^L(\mathcal{I}^L(t)) = \sum g_i^*(\mathcal{I}^L(t,\tau), \mathcal{I}^F(t,\tau)) P(\mathcal{I}^F(t,\tau)|\mathcal{I}^L(t)),$$

where the sum is over all $\mathcal{I}^F(t,\tau)$. Further, g_i^* is the unique solution of the equation

$$g_i^*(\mathcal{I}^L(t,\tau), \mathcal{I}^F(t,\tau)) = \sum\nolimits^1 P(s(t), a(t)|\mathcal{I}^F(t,\tau), \mathcal{I}^L(t,\tau))\{c_i^L(s(t), a(t))$$
$$+ \beta \sum\nolimits^2 g_i^*[(\mathfrak{z}^L(t+1), \mathcal{I}^L(t,\tau-1)), (\mathfrak{z}^F(t+1), \mathcal{I}^F(t,\tau-1))] P(z(t+1), s(t+1)|s(t), a(t))\},$$

where $\mathfrak{z}^k(t) = \{z^k(t), s^k(t), a^k(t-1)\}$, $\sum^1$ *is over all* $s(t)$ *and* $a(t)$, *and* $\sum^2$ *is over all* $z(t+1)$ *and* $s(t+1)$.

Proof. We remark that since (π^L, π^F) is assumed given, $P(s(t), a(t)|\mathcal{I}^F(t,\tau), \mathcal{I}^L(t,\tau))$ is well defined. Assume there is a function g such that

$$v(\mathcal{I}^L(t)) = \sum g(\mathcal{I}^L(t,\tau), \mathcal{I}^F(t,\tau)) P(\mathcal{I}^F(t,\tau)|\mathcal{I}^L(t)),$$

where the sum is over all $\mathcal{I}^F(t,\tau)$. Then it is straightforward to show that there is a function g' such that

$$h_i^L(\mathcal{I}^L(t), v) = \sum g'(\mathcal{I}^L(t,\tau), \mathcal{I}^F(t,\tau)) P(\mathcal{I}^F(t,\tau)|\mathcal{I}^L(t)),$$

where the sum is over all $\mathcal{I}^F(t,\tau)$, and

$$g'(\mathcal{I}^L(t,\tau), \mathcal{I}^F(t,\tau)) = \sum{}^1 P(s(t), a(t)|\mathcal{I}^F(t,\tau), \mathcal{I}^L(t,\tau)) \{c_i^L(s(t), a(t))$$
$$+ \beta \sum{}^2 g[(\mathfrak{z}^L(t+1), \mathcal{I}^L(t,\tau-1)), (\mathfrak{z}^F(t+1), \mathcal{I}^F(t,\tau-1))] P(z(t+1), s(t+1)|s(t), a(t))\},$$

and where $\mathfrak{z}^k(t) = \{z^k(t), s^k(t), a^k(t-1)\}$, $\sum^1$ is over all $s(t)$ and $a(t)$, and $\sum^2$ is over all $z(t+1)$ and $s(t+1)$.The result follows directly from the following facts:

- The operator H^L, where $[H^L v](\mathcal{I}^L(t)) = h_i^L(\mathcal{I}^L(t), v)$, is a contraction operator on the Banach space comprised of all functions mapping $\mathcal{I}^L(t)$ into the real line, having as its norm the supremum norm.
- As a result, the sequence $\{v^n\}$, where $v^{n+1} = H^L v^n$, converges to v^L for any given v^0.

□

Since both π^L and π^F are finite-memory policies, then both $\mathcal{I}^L(t,\tau)$ and $y^L(t)$ are t-invariant arrays of finite dimension, which enhances the potential computability of

v^L.We remark that Proposition 2 holds for any given finite memory leader policy ρ^L and follower policy ρ^F, and hence ρ^F is not necessarily a response policy to ρ^L.

We now summarize how fitness measures are determined for a given finite-memory leader policy:

- Step 1: Determine a *perfect* memory follower response policy that achieves the minimum expected cost for the follower, using Proposition 1.
- Step 2: Approximate the resulting perfect-memory follower response policy by a *finite*-memory policy.
- Step 3: Given the leader policy and the follower's approximate response policy, determine the concomitant fitness measures using Proposition 2.

3.5 Multi-Objective Genetic Algorithm

We now describe how we use a multi-objective genetic algorithm (MOGA), NSGA-II (Deb, 2002), to generate policies from which the leader will choose a most preferred policy.

Let $\{\pi^L(m), m = 1, \ldots, M\}$ be the current population of the leader's finite memory policies, and for each m, let $\pi^F(m)$ be the finite memory follower's response policy to the leader's policy $\pi^L(m)$. Further, let $v_i^L(\pi^L(m), \pi^F(m))$ be the expected cost of the leader's i^{th} criterion, given the policy pair $(\pi^L(m), \pi^F(m))$.

Definition Policy π^L is said to *dominate* policy ρ^L if

$$v_i^L(\pi^L, \pi^F) \leq v_i^L(\rho^L, \rho^F), \forall i$$

and there exists at least one i such that

$$v_i^L(\pi^L, \pi^F) < v_i^L(\rho^L, \rho^F),$$

where π^F and ρ^F are the follower's response policies to the leader's policy π^L and ρ^L, respectively.

Policy π^L is said to be *non-dominated* if there does not exist a policy that dominated policy π^L.

The MOGA constructs the next generation of policies from the current set of policies as follows. The MOGA encodes each policy $\pi^L(m)$ into a chromosome. A gene is an element of the chromosome vector, and an allele is a numerical value taken by a gene. In the context of our model, the chromosome is a probability mass vector over the action space, and the i^{th} gene of the chromosome denotes the probability that action i is selected by the leader. The MOGA then determines $v_i^L, 1 \leq i \leq N$ for each chromosome, where N is the number of leader's objectives. A description of how v_i^L is determined can be found in section 3.4. The tuple $(v_1^L, ..., v_N^L)$ serves as the fitness measure of this chromosome.

On the basis of $(v_1^L, ..., v_N^L)$, the population of chromosomes are partitioned into subsets called fronts, where front 1 is the set of non-dominated chromosomes, and front $k+1$ is the set of non-dominated chromosomes when the chromosomes in fronts 1 through k are removed from consideration, $k = 1, 2, \ldots$. Chromosomes in front k are given rank k. In addition, the crowding distance of each chromosome is determined within each front. Crowding distance is defined as the average Euclidean distance of a chromosome to the other chromosomes in the front, based on $(v_1^L, ..., v_N^L)$ as a measure of position. Crowding distance is considered a measure of diversity for the policies, and for this measure, larger is considered better. The current generation of policies is sorted according to ranking and crowding distance.

Parents are selected from the current generation of policies, based on their ranks and crowding distances. Chromosomes with higher rank and larger crowding distance are selected to generate offspring with higher probability. The selected parents form a mating pool and generate offspring using a crossover operator. For each parents pair, the

crossover operator randomly exchanges a portion of genes with each other to form two new offspring. A mutation operator is also used to maintain genetic diversity from one generation to another. This operator randomly alters a certain percentage of genes in the current generation of policies. Then the non-dominated sorting procedure is applied again on the current population and offspring population, the top M (population size) chromosomes are kept and this is the next population. The whole algorithm repeats for a certain number of iterations.

3.6 Equilibria

We remark that there are two equilibrium conditions, one associated with each agent. With respect to the follower, assume $\bar{\pi}^F$ is the *perfect* memory response policy to a given leader policy π^L. Then, results in Proposition 1 imply that

$$v^F(\mathcal{I}^F(t)) = v^F(\pi^L, \bar{\pi}^F; \mathcal{I}^F(t)) \leq v^F(\pi^L, \rho^F; \mathcal{I}^F(t))$$

for all follower policies ρ^F and all $\mathcal{I}^F(t)$, where a direct application of the results of Proposition 2 can be used to determine $v^F(\rho^L, \rho^F; \mathcal{I}^F(t))$ for any pair of leader-follower policies (ρ^L, ρ^F).

With respect to the leader, we now assume π^F is a *finite-memory approximation* of the perfect memory follower's response policy to the given leader policy π^L. Let $v^L(\pi^L, \pi^F; \mathcal{I}^L(t))$ be the vector of criterion values for the leader's multiple objectives, given (π^L, π^F), and information state $\mathcal{I}^L(t)$. Our process of determining candidate leader policies from which the leader can select the most preferred policy is intended to determine (π^L, π^F) pairs so that there exists no pair (ρ^L, ρ^F) such that

$$v_i^L(\rho^L, \rho^F; \mathcal{I}^L(t)) \leq v_i^L(\pi^L, \pi^F; \mathcal{I}^L(t)), \forall i$$

for all $\mathcal{I}^L(t)$, where ρ^L represents any leader policy and ρ^F represents a finite-memory

approximation of the perfect memory follower response policy to ρ^L.

We note that by Proposition 1 and results in (Puterman, 1994, Chapter 6), the follower policy in the first equilibrium condition is an optimal policy; hence, the follower has no incentive to deviate from this policy. With respect to the second equilibrium condition, we note that all of the follower's policies are finite memory approximations of the follower's optimal perfect memory response policy to the leader's policy. Further, the process of determining the leader policies does not guarantee that pairs (π^L, π^F) will be determined that satisfy the second equilibrium condition. Thus, there is no guarantee that the leader will not want to deviate from the set of resultant non-dominated leader policies. However, given a sufficient number of generations of the MOGA and a sufficiently large τ such that the finite memory follower policy is a good approximation to an optimal perfect memory follower policy, it is reasonable to be confident that the resultant leader will have little incentive to deviate from the set of non-dominated leader policies generated. Figure 1 provides an outline of the process for generating these non-dominated leader policies.

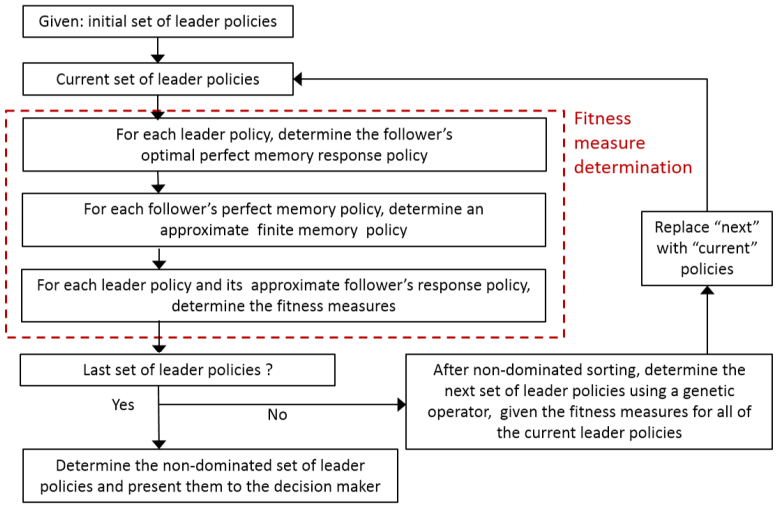
Given: initial set of leader policies
Current set of leader policies
For each leader policy, determine the follower's optimal perfect memory response policy
For each follower's perfect memory policy, determine an approximate finite memory policy
For each leader policy and its approximate follower's response policy, determine the fitness measures
Fitness measure determination
Last set of leader policies ?
Yes
No
Determine the non-dominated set of leader policies and present them to the decision maker
After non-dominated sorting, determine the next set of leader policies using a genetic operator, given the fitness measures for all of the current leader policies
Replace "next" with "current" policies

Figure 1: Outline of the decision support process

3.7 Value of Information

We now address the question: will improved observation quality improve, or at least not degrade, the performance of agents? With respect to the follower, assume $P(z(t+1)|s(t+1), s(t), a(t)) = P(z^F(t+1)|s^L(t+1), a(t))P(z^L(t+1)|s^F(t+1), a(t))$, and let $P^F(a(t))$ be the stochastic matrix having ij^{th} entry $P(z^F(t+1)|s^L(t+1), a(t))$, where $i = s^L(t+1)$ and $j = z^F(t+1)$. Let $Q^F(a(t))$ be a stochastic matrix such that for all $a(t)$, there exists a third stochastic matrix $R^F(a(t))$, where $Q^F(a(t)) = R^F(a(t))P^F(a(t))$. Then by results in (White and Harrington, 1980), $v^P(\mathcal{I}^F(t)) \leq v^Q(\mathcal{I}^F(t))$, where $v^P(\mathcal{I}^F(t))$ is the value of the follower's criterion associated with observation matrix $P^F(a(t))$ and $v^Q(\mathcal{I}^F(t))$ is the value of the follower's criterion associated with observation matrix $Q^F(a(t))$. Thus, for the follower, the observation quality provided by $P^F(a(t))$ is at least as good as the observation quality provided by $Q^F(a(t))$ and the added value of using $P^F(a(t))$, relative to $Q^F(a(t))$, is the difference $0 \leq v^Q(\mathcal{I}^F(t))$–$v^P(\mathcal{I}^F(t))$.

The determination of conditions that guarantee that improved observation quality for the leader will not degrade leader performance is a topic for future research. We remark that counterexamples exist to the general claim that better information quality always implies improved system performance (Ortiz, Erera and White, 2012).

4 Numerical Results

This section presents an example that illustrates how the results in Section 3 can be used to support a defender (the leader) in selecting a sequence of actions to best control a simplified liquid egg production process over time, given that there is an attacker (the follower) who seeks to contaminate the liquid egg production process with a chemical or biological toxin. Both the defender and the attacker receive updated, possibly noise corrupted, data about his/her opponent just prior to each decision epoch. The defender

has two scalar criteria: (1) a measure of the system's vulnerability to an attacker, which the defender wants to minimize, and (2) a measure of the system's productivity, which the defender wants to maximize. The attacker's criterion is to maximize the expected number of packages produced by the facility that contain a sufficiently lethal dose of the toxin.

The attacker's and defender's transition diagrams are presented in Figures 2 and 3, respectively. Two targets, T_1 and T_2 are considered. State O is the state that the attacker is in prior to launching an attack, and includes attack team assembly, toxin manufacture, and transportation. States PT_1 and PT_2 are the pre-attack states in which the attacker is armed and ready to attack target 1 and target 2, respectively. At each decision epoch, the attacker can either choose to stay in the current state, advance forward to the next state, or retreat to a prior state. We remark that the transition probability of the attacker could be affected by the defender's strategy. An attacker's error or interdiction by the defender could return the attacker back to the state O.

The defender's states include a full production low alert state (FP), a low production high alert state (LP) and the attacked state (Att.). The defender in FP or LP could stay in his/her current state or switch to the other state with given probabilities. If an attack occurs, the failure of detecting an attack promptly and remain in either FP or LP will cause significant consequences for the defender. The defender can also terminate the game and shut down the facilities to clean the toxin if the attack is successfully detected.

Given an attack, there are two groups of outcomes: unsuccessful attacks and successful attacks. If the attack is successful, the amount of toxin delivered to the consumers is described by a cumulative distribution function. We assume that the distribution function satisfies a standard, first-order stochastic dominance assumption.

We restrict our attention to deterministic defender policies since they are easy to implement. Related discussion can be found in Paruchuri (2004). There are 64 defender policies and the MOGA can probabilistically identify the Pareto efficient policies within

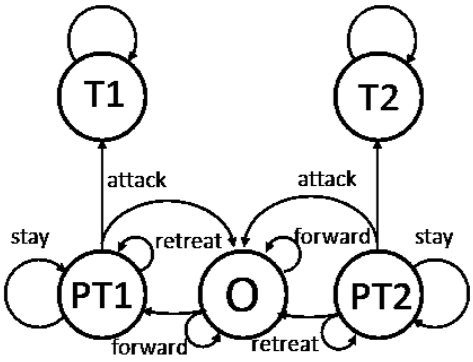

Figure 2: Attacker's transition diagram

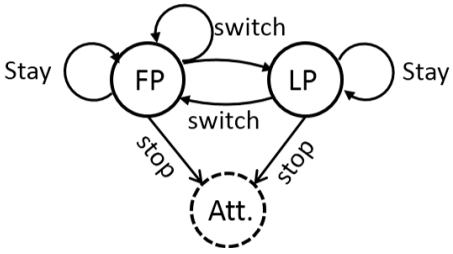

Figure 3: Defender's transition diagram

5 generations. The non-dominated set of policies are presented in Table 1. The defender must then trade off productivity and vulnerability, based on his/her preferences in order to select a most preferred policy.

Table 1: Decision Support Table

Policy	Productivity ratio to maximum	Vulnerability ratio to minimum
π_1	1.000	7.077
π_2	0.959	5.737
π_3	0.942	5.633
π_4	0.847	4.136
π_5	0.691	1.000

Figure 4 compares the value of the attacker's reward, as a function of observation

quality. The results are consistent with the discussion in Section 3.7; i.e. improved observation quality of the follower does not degrade the follower's performance. A comparison of the value of the defender's rewards as a function of observation quality is a topic of future research.

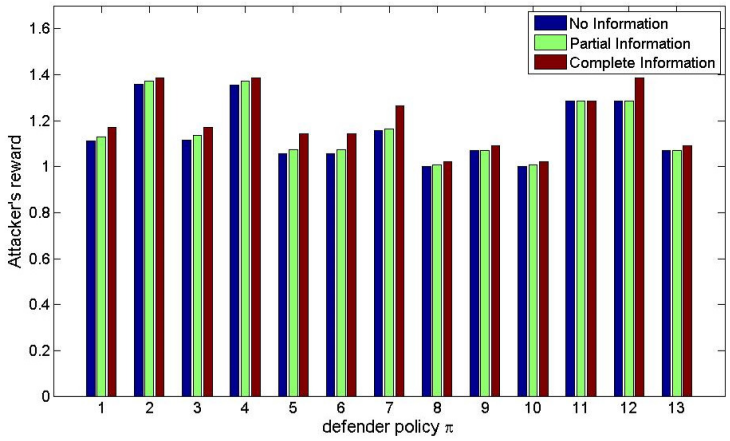

Figure 4: Attacker's reward under different qualities of observations.

5 Conclusions and Extensions

The contributions of the paper are as follows:

(1) We have blended the POMG, the POMDP, and the MOGA to identify leader policies that are candidates for a most preferred policy in an infinite horizon, sequential decision making environment where:

- There are two intelligent and adaptable agents, a leader and a follower, and each can affect the performance of the other.
- At each decision epoch, each agent knows: its past and present states, its past

actions, and the noise corrupted observations of the other agent's past and present states.

- The leader's and follower's actions are selected simultaneously at each decision epoch.
- Each agent's policy selects actions based on data currently available to the agent.
- The follower knows the leader's policy and determines a response policy that is optimal with respect to the follower's objective.
- The leader considers multiple objectives in selecting its policy.

(2) Given the POMG, a leader policy, and the assumption that the follower selects its policy with complete knowledge of and in response to the policy selected by the leader, we have constructed a specially structured POMDP that leads to the determination of a perfect-memory optimal policy for the follower (Proposition 1). We have shown that this POMDP has a computationally useful sufficient statistic and a value function structure described in terms of this sufficient statistic. By assuming that the leader policy is a finite-memory policy, we have shown that the sufficient statistic is finite-dimensional and that the value function has at least an approximate finite representation, thus insuring that at least a near-optimal perfect-memory policy for the follower is potentially computable.

(3) We have determined a computationally tractable procedure for calculating the fitness measures for the MOGA, given that the policies for both agents are finite-memory policies (Proposition 2). We have presented a simple procedure for finding a finite-memory approximation to a perfect-memory policy and used it to find a finite-memory policy for the follower, based on the perfect-memory policy determined through the use of Proposition 1. The concomitant results show that there is a finite-dimensional sufficient statistic for the related fitness measures.

We remark that the computational tractability of procedures for determining the follower's response policy and the fitness measures for the MOGA is inextricably linked to the assumption that the agent policies are finite-memory policies.

The output of the process described in this paper can serve as the options generation phase of an option selection process; e.g., a deterministic version of multi-attribute decision theory (Kenney & Raiffa, 1993).

Topics for future research include a sensitivity analysis in order to better understand the robustness of the results and a study of the value of improving the leader's quality of observations.

Acknowledgments

This material is based upon work supported by the U.S. Department of Homeland Security under Grant Award Number 2010-ST-061-FD0001 through a grant awarded by the National Center for Food Protection and Defense at the University of Minnesota. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security or the National Center for Food Protection and Defense.

References

[1] Aberdeen, D., A (revised) survey of approximate methods for solving partially observable Markov decision processes, *Technical report*, Research School of Information Science and Engineering, Australia National University, 2003.

[2] Bakir, N. O., A Stackelberg game model for resource allocation in cargo container security, *Annals of Operations Research*, **187**, 5 - 22, 2011.

[3] Bakir, N. O. and Kardes, K., A Stochastic game model on overseas cargo container security, *Non-published Research Reports*, CREATE center, Paper 6, 2009.

[4] Becker, R., Zilberstein, S., Lesser, V., and Goldman, C. V., Solving transition independent decentralized Markov decision processes, *Journal of Artificial Intelligence Research (JAIR)*, **22**, 423 - 455, 2004.

[5] Bernstein, D. S., Givan, R., Immerman, N., and Zilberstein, S., The complexity of decentralized control of Markov decision processes, *Mathematics of Operations Research*, **27**(4), 819 - 840, 2002.

[6] Bernstein, D. S., Hansen, E. A., and Zilberstein, S., Bounded policy iteration for decentralized POMDPs, In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI)*, 1287 - 1292, Edinburgh, Scotland, July 2005.

[7] Bier, V. M., Oliveros, S., and Samuelson, L. Choosing what to protect: Strategic defensive allocation against an unknown attacker, *Journal of Public Economic Theory*, **9**(4), 563 - 587, 2007.

[8] Bier, V. M., Haphuriwat, N., Menoyo, J., Zimmerman, R., and Culpen, A. M., Optimal resource allocation for defense of targets based on differing measures of attractiveness, In *Risk Analysis*, **28**(3), 763 - 770, 2008.

[9] Bopardikar, S. D., and Hespanha, J. P., Randomized solutions to partial information dynamic zero-sum games, In *American Control Conference (ACC)*, San Francisco, CA, June 2011.

[10] Bowman, M., Briand, L. C., and Labiche, Y., Solving the class responsibility assignment problem in object-oriented analysis with multi-objective genetic algorithms, *IEEE Transactions on Software Engineering (TSE)*, **36**(6), 817 - 837, 2010.

[11] Cardoso,J.M.P, Diniz, P.C , Game Theory Models of Intelligent Actors in Reliability Analysis: An Overview of the State of the Art, *Game Theoretic Risk Analysis of Security Threats, International Series in Operations Research & Management Science*, **128**, 1-19, 2009.

[12] Cassandra, A. R., Kaelbling, L. P., and Littman, M. L., Acting optimally in partially observable stochastic domains , in *Proceedings Twelfth National Conference on Artificial Intelligence (AAAI-94)*, Seattle, WA, 1023 - 1028, 1994a.

[13] Cassandra, A. R., Optimal policies for partially observable Markov decision processes, *Technical Report (CS-94-14)*, Brown University, Department of Computer Science, Providence RI, 1994b.

[14] Cassandra, A. R., Littman, M. L. and Zhang, N. L., Incremental pruning: a simple, fast, exact method for partially observable Markov decision processes, in *Proceedings Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97)*, Morgan Kaufmann, San Francisco, CA, 54 - 61, 1997.

[15] Cavusoglu, H., and Kwark, Y., Passenger profiling and screening for aviation security in the presence of strategic attackers, *Decision Analysis*, **10**(1), 63 - 81, 2013.

[16] Cheng, H. T., Algorithms for partially observable Markov decision processes, PhD thesis, University of British Columbia, Vancouver, British Columbia, Canada, 1988.

[17] Coello, C. A. C., An updated survey of GA-Based multiobjective optimization techniques, *ACM Computing Survey*, **32**(2), 109 - 143, 2000.

[18] Corne, D. W., Knowles, J. D. and Oates, M. J., The Pareto envelope-based selection algorithm for Multiobjective optimization, In *Proceedings of the Parallel Problem Solving from Nature VI Conference*, **1917**, 839 - 848, 2000.

[19] Deb, K., Nonlinear goal programming using multi-objective genetic algorithms, *Journal of the Operational Research Society*, **52**(3), 291 - 302, 2001.

[20] Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Transactions on Evolutionary Computation*, **6**(2), 182 - 197, 2002.

[21] Eagle, J. N., The optimal search for a moving target when the search path is constrained, *Operations Research*, **32**(5), 1107 - 1115, 1984.

[22] Emery-Montemerlo, R., Gordon, G., Schneider, J., and Thrun, S., Approximate solutions for partially observable stochastic games with common payoffs, In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 136 - 143, 2004.

[23] Feng, Z., and Zilberstein, S., Region-based incremental pruning for POMDPs, In *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence (UAI-04)* , Morgan Kaufmann, San Francisco, 146 - 153, 2004.

[24] Filar, J., and Vrieze, K., Competitive Markov decision processes, Springer, Heidelberg, 1997.

[25] Forrest, S., Genetic algorithms: principles of natural selection applied to computation, *Science*, **261**, 872 - 878, 1993.

[26] Ghosh, M. K., McDonald, D., and Sinha, S., Zero-sum stochastic games with partial information, *Journal of optimization theory and applications*, **121**(1), 99 - 118, 2004.

[27] Gmytrasiewicz, P. J., and Doshi, P., A framework for sequential planning in multiagent settings, *Journal of Artificial Intelligence Research*, **24**, 49 - 79, 2005.

[28] Goldberg, D. E., Genetic algorithms in search, optimization, and machine learning, Addison-Wesley: Reading, MA, 1989.

[29] Hansen, E. A., An improved policy iteration algorithm for partially observable MDPs, *Advances in Neural Inform. Processing Systems*, **10** (NIPS-97), 1015 - 1021, MIT Press, Cambridge, MA, 1998a.

[30] Hansen, E. A., Solving POMDPs by searching in policy space, in *Proceedings of Uncertainty in Artificial Intelligence*, **10**, 211 - 219, 1998b.

[31] Hansen, E. A., Bernstein, D. S., and Zilberstein, S., Dynamic programming for partially observable stochastic games, in *Proceedings of the nineteenth national conference on artificial intelligence*, 709 - 715, San Jose, California, 2004.

[32] Hausken, K., and Zhuang, J., Governments' and terrorists' defense and attack in a T-period game, *Decision Analysis*, **8**(1), 46 - 70, 2011.

[33] Hauskrecht M., Value-function approximations for partially observable Markov decision processes, *Journal of Artificial Intelligence Research*, **13**, 33 - 94, 2000.

[34] Hauskrecht M., Planning and control in stochastic domains with imperfect information, PhD thesis, Massachusetts Institute of Technology, 1997.

[35] Hespanha, J. P., and Prandini, M., Nash equilibria in partial-information games on Markov chains, in *IEEE Conference on Decision and Control*, Orlando, FL, 2102 - 2107, Dec 2001.

[36] Hoey, J., Poupart, P., von Bertoldi, A., Craig, T., Boutilier, C., and Mihailidis, A., Automated handwashing assistance for persons with dementia using video and a partially observable Markov decision process, *Computer Vision and Image Understanding*, **114** (5), 503 - 519, 2010.

[37] Holland, J. H., Adaptation in natural and artificial systems, University of Michigan Press: Ann Arbor, 1975. Reprinted in 1992 by MIT Press, Cambridge MA.

[38] Holloway, H, and White, C. C., Question Selection and Resolvability for Imprecise Multi-Attribute Alternative Selection, *IEEE Transactions on Systems, Man, and Cybernetics*, Part A, **38**(1),162-169, 2008.

[39] Horn, J., Nafpliotis, N., and Goldberg, D. E., A niched Pareto genetic algorithm for multiobjective optimization, In *Proceedings of the 1st IEEE conference on evolutionary computation, IEEE World Congress on Computational Intelligence* **1**, 82 - 87, Orlando, FL, June, 1994.

[40] Keeney, R. L. and Raiffa, H., Decisions with multiple objectives: preferences and value trade-offs, Cambridge University Press, 1993.

[41] Koller, D., Meggido, N., and von Stengel, B., Fast algorithms for finding randomized strategies in game trees, In *26th Annual ACM Symposium on the Theory of Computing*, 750 - 759, 1994.

[42] Konak, A., Coit, D. W., and Smith, A. E., Multi-objective optimization using genetic algorithms: A tutorial, *Reliability Engineering and System Safety*, **91**, 992 - 1007, 2006.

[43] Kumar, A., and Zilberstein, S., Dynamic programming approximations for partially observable stochastic games, In *Proceedings of the twenty-second international FLAIRS conference*, 547 - 552, Sanibel Island, Florida.

[44] Lin, A. Z.-Z.,Bean, J., and White, C. C., Genetic algorithm heuristics for finite horizon partially observed Markov decision problems, *Technical Report*, University of Michigan, Ann Arbor, 1998.

[45] Lin, A. Z.-Z.,Bean, J., and White, C. C., A hybrid genetic/optimization algorithm for finite horizon partially observed Markov decision processes, *Journal on Computing*, **16**(1), 27 - 38, 2004.

[46] Lin, C. M. and Gen, M., Multi-criteria human resource allocation for solving multistage combinatorial optimization problems using multiobjective hybrid genetic algorithm, *Expert Systems with Applications*, **34**(4), 2480 - 2490, 2008.

[47] Littman, M. L., Memoryless policies: theoretical limitations and practical results, In *Proceedings of the Third International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, 238 - 245, 1994a.

[48] Littman, M. L., The Witness algorithm: solving partially observable Markov decision processes, Brown University, Department of Computer Science, *Technical Report* CS-94-40, 1994b.

[49] Lovejoy, W. S., A survey of algorithmic methods for partially observed Markov decision process, *Annals of Operations Research*, **28** (1), 47 - 65, 1991a.

[50] Manning, L., Baines, R., and Chadd, S., Deliberate contamination of the food supply chain, *British Food Journal*, **107** (4), 225 - 245, 2005.

[51] McEneaney, W. M., Some classes of imperfect information finite state-space stochastic games with finite-dimensional solutions, *Applied Mathematics and Optimization*, **50**(2), 87 - 118, 2004.

[52] McLay, L.,Rothschild, C., Guikema, S., Robust adversarial risk analysis: A level-k approach *Decision Analysis*, **9**(1), 41 - 54, 2012.

[53] Monahan, G. E., A survey of partially observable Markov decision processes: Theory, models, and algorithms, *Management Science*, **28**, 1 - 16, 1982.

[54] Mohtadi, H. and Murshid, A. P., Risk analysis of chemical, biological, or radionuclear threats: Implications for food security, *Risk Analysis*, **29**, 1317 - 1335, 2009.

[55] Naser-Moghadasi, M., Evaluating effects of two alternative filters for the incremental pruning algorithm on quality of POMDP exact solutions, *International Journal of Intelligence Science*, **2**, 1 - 8, 2012.

[56] Oliehoek, F. A., Spaan, M. T. J., and Vlassis, N., Best-response play in partially observable card game, In *Proceedings of the 14th Annual Machine Learning Conference of Belgium and the Netherlands*, 45 - 50, Feb. 2005.

[57] Oliehoek, F. A., Spaan, M. T. J., and Vlassis, Nikos, Optimal and approximate Q-value functions for decentralized POMDPs, *Journal of Artificial Intelligence Research*, **32**, 289 - 353, 2008.

[58] Oliehoek, F. A.,Decentralized POMDPs, In: M. Wiering, & M. V. Otterlo (Eds.), *Reinforcement learning: State of the art*, **12**, 471 - 503, Springer, 2012.

[59] Ombuki, B., Ross, B. J., and Hanshar, F., Multi-objective genetic algorithms for vehicle routing problem with time windows, *Applied Intelligence*, **24**, 17 - 30, 2006.

[60] Ortiz, O. L., Erera, A. L., White, C. C., State observation accuracy and finite-memory policy performance for partially observed Markov decision processes, *technical report*, 2012.

[61] O'Ryan, M., Djuretic, T., Wall, P., Nichols, G., Hennessy, T., Slutsker, L., Hedberg, C., MacDonald, K., and Osterholm, M., An outbreak of salmonella infection from ice cream, *New England Journal of Medicine*, **335**(11), 824 - 825, 1996.

[62] Paruchuri, P., Tambe, M., Ordonez, F., and Kraus, S., Towards a formalization of teamwork with resource constraints, In *International Joint Conference on Autonomous Agents and Multiagent Systems*, 596 - 603, 2004.

[63] Pineau, J., Gordon, G. J., and Thrun, S., Point-based value iteration: An anytime algorithm for POMDPs, In *International Joint Conference on Artificial Intelligence*, 1025 - 1032, 2003.

[64] Platzman, L. K., Finite memory estimation and control of finite probabilistic systems, PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1977.

[65] Platzman, L. K., Optimal infinite-horizon undiscounted control of finite probabilistic systems, *SIAM J. Control Optim.*, **18**, 362 - 380, 1980.

[66] Ponnambalam, S. G., Ramkumar, V., and Jawahar, N., A multiobjective genetic algorithm for job shop scheduling, *Production Planning & Control: The Management of Operations*, **12**(8), 764 - 774, 2001.

[67] Poupart, P., Exploiting structure to efficiently solve large scale partially observable Markov decision processes, PhD thesis, Department of Computer Science, University of Toronto, 2005.

[68] Poupart, P., Kim, K. E., and Kim, D., Closing the gap: Improved bounds on optimal POMDP solutions, In *International conference on planning and scheduling (ICAPS)*, 2011.

[69] Puterman, M. L., Markov decision processes: discrete dynamic programming, New York: J Wiley & Sons, 1994.

[70] Rabinovich, Z., Goldman, C. V., and Rosenschein, J. S., The complexity of multi-agent systems: the price of silence, In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1102 - 1103, Melbourne, Australia, 2003.

[71] Raghavan, T. E. S., and Filar, J. A., Algorithms for stochastic games - a survey, *Methods and Models of Operations Research*, **35**, 437 - 472, 1991.

[72] Rothschild, C., McLay, L., Guikema, S., Adversarial risk analysis with incomplete information:A level-k approach, *Risk Analysis*, **32**(7), 1219 - 1231, 2012.

[73] Schaffer, J. D., Multiple objective optimisation with vector evaluated genetic algorithm, In *Proceedings of the 1st International Conference on Genetic Algorithms*, 93 - 100, Morgan Kaufmann Publishers, Inc., San Mateo, 1985.

[74] Seuken, S., and Zilberstein, S., Improved memory-bounded dynamic programming for decentralized POMDPs, In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, Vancouver, Canada, July, 2007.

[75] Shani, G., Pineau, J., and Kaplow, R., A survey of point-based POMDP solvers, *AAMAS*, 1 - 51, June 2012.

[76] Shapley, L. S., Stochastic games, *Proceedings of the National Academy of Sciences of the U. S. A.*, **39**, 1095 - 1100, 1953.

[77] Smallwood, R. D., and Sondik, E. J., The optimal control of partially observable Markov decision processes over a finite horizon, *Operations Research*, **21**, 1071 - 1088, 1973.

[78] Sobel, J., Khan, A., and Swerdlow, D., Threat of a biological terrorist attack on the US food supply: the CDC perspective, *The Lancet*, **359**(9309), 874 - 880, 2002.

[79] Sondik, E. J., The optimal control of partially observable Markov processes over the infinite horizon: discounted costs, *Operations Research*, **26**, 282 - 304, 1978.

[80] Srinivas, M., and Patnaik, L. M., Genetic algorithms: A survey, *IEEE Computer*, **27**(6), 17 - 26, 1994.

[81] Thomson, B., and Young, S., Bayesian update of dialogue state: A POMDP framework for spoken dialogue systems, *Computer Speech & Language*, **24**(4), 562 - 588, 2010.

[82] Ummels, M., Stochastic multiplayer games: theory and algorithms, PhD thesis, RWTH Aachen University, 2010.

[83] Varakantham, P., Maheswaran, R., Gupta, T., and Tambe, M., Towards efficient computation of quality bounded solutions in POMDPs: Expected Value Approximation and Dynamic Disjunctive Beliefs. In *Twentieth International Joint Conference on Artificial Intelligence*, 2007.

[84] Wang, C. and Bier, V. M., Target-hardening decisions based on uncertain multiattribute terrorist utility, *Decision Analysis*, **8**(4), 286 - 302, 2011.

[85] White, C. C., and Harrington, D. P., Application of Jansen's inequality to adaptive suboptimal design, *Journal of Optimization Theory and Application*, **32**, 89 - 99, 1980.

[86] White, C. C., A survey of solution techniques for the partially observed Markov decision process, *Annals of Operations Research*, **32**, 215 - 230, 1991.

[87] White, C. C., and Scherer, W. T., Solution procedures for partially observed Markov decision processes, *Operations Research*, **37**, 791 - 797, 1989.

[88] White, C. C., and Scherer, W. T., Finite-memory suboptimal design for partially observed Markov decision processes, *Operations Research*, **42**, 439 - 455, 1994.

[89] Yildirim, M. B., and Mouzon, G., Single-machine sustainable production planning to minimize total energy consumption and total completion time using a multiple objective genetic algorithm, *IEEE Transactions on Engineering Management*, **59**(4), 585-597, 2012.

[90] Yu, H., Approximation solution methods for partially observable Markov and semi-Markov decision processes, PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 2007.

[91] Zhang, H., Partially observable Markov decision processes: a geometric technique and analysis, *Operations Research*, **58**, 214 - 228, 2010.

[92] Zhuang J. and Bier, V. M., Balancing terrorism and natural disasters defensive strategy with endogenous attacker effort, *Operations Research*, **55**(5), 976 - 991, 2007.